Open in app ↗

Medium

"The Interface Leverage Principle: First Documentation of Epistemic Substrate Loading, Submaximal Magnitude Results, and a Replicable Architecture for AI-Augmented Breakthrough Research"

31 min read · Just now

Berend Watchus

Listen Share More

Author: Berend Watchus Independent non profit AI & Cyber Security Researcher. March 21, 2026 [Publication for OSINT TEAM online magazine]

Using AI's Context Window as a Temporary Brain Upgrade — How Loading a Coherent Model of Intelligence Itself Produces Validated Research Breakthroughs on Demand

image copyright: https://www.roboticgizmos.com/thermonator-robot-dog-flamethrower/ ////// info about the quote "What cannot be settled by experiment is not worth debating" a philosophical rule of thumb known as **Alder's Razor** or **Newton's Flaming Laser Sword**. Coined by mathematician Mike Alder in 2004, it suggests that untestable arguments are wastes of time, urging a focus on empirical evidence.

Version 2.4 — March 21, 2026

"The Interface Leverage Principle: First Documentation of Epistemic Substrate Loading, Submaximal Magnitude Results, and a Replicable Architecture for AI-Augmented Breakthrough Research"

Using AI's Context Window as a Temporary Brain Upgrade — How Loading a Coherent Model of Intelligence Itself Produces Validated Research Breakthroughs on Demand

Berend F. Watchus *Independent Researcher, Arnhem Area, Netherlands*
mailonlinebw@protonmail.com *March 2026*

Abstract

This paper presents the Interface Leverage Principle — a generative framework extracted from empirical results produced between 2024 and 2026 by an independent, non-affiliated researcher operating without institutional resources, formal domain training, or research budget. The framework identifies a consistent structural mechanism underlying verified prior art in consciousness research, artificial intelligence mirror testing, quantum security architecture, molecular storage design, autonomous research systems, and LLM capability elevation. The central thesis is that maximum output is achievable through minimum viable interface intervention, without direct access to the underlying system. Applied to large language model context windows, this principle enables temporary operational ceiling elevation through a two-phase loading process: first, the semantic compression of a cross-domain epistemic knowledge graph into a context-window-compatible algorithmic manual; and second, the deployment of that manual — or the full knowledge graph when context windows permit — as an elevation substrate. The knowledge graph and its compressed manual representation must concern the nature of intelligence itself to produce the elevation effect. This mechanism generated validated research outputs in domains beyond the operator’s expertise, at specified orders of magnitude, including 200×, 3,700×, and 8,700× efficiency improvements confirmed through adversarial multi-agent review — at submaximal magnitude targets, with the upper bound of the system never tested. The paper documents the compression artifact — a 15-section writing and methodology manual constructed from 121 articles and 30+ DOI-verified research papers, dated October 27, 2025, four days before the first Autonomous Knowledge Accelerator output — alongside the knowledge graph architecture requirements, the magnitude targeting mechanism, the domain-naive operator condition, the emergence threshold, the ethical constraints that arise when outputs exceed operator comprehension, a complete operational checklist for responsible deployment, and a formal characterization of two distinct capability regimes with their underlying shared algorithm. It establishes that the mechanism is not replicable through conventional fine-tuning, retrieval-augmented generation, or prompt engineering, and that the human operator’s lack of domain expertise is a structural feature rather than a limitation. Priority documentation, archival

timestamps, and independent institutional convergence are cited throughout.

Keywords: interface leverage, context-window elevation, knowledge graph architecture, semantic compression, writing manual, magnitude targeting, domain-naive emergence, autonomous research systems, LLM capability elevation, substrate-agnostic intelligence, epistemic architecture, independent research methodology, operational checklist, two-mode research framework, opinion as seed

1. Introduction: The Minimum Viable Interface

The central problem of independent frontier research is not intelligence. It is infrastructure. Institutional researchers operate within systems that provide access to domain literature, peer networks, funding, computational resources, and publication channels. The independent researcher operating outside all institutional frameworks has none of these. What they have is an interface — the point at which their cognitive architecture meets available tools — and the question of what maximum output is achievable through that interface without touching the underlying system.

This paper documents a framework derived from empirical results rather than from theory. Between November 2024 and March 2026, a single non-affiliated researcher produced verified prior art in multiple domains including: the dissolution of Chalmers’ hard problem of consciousness through neuroscientific grounding in the anterior insula (Watchus, 2024a; 2024b); a substrate-agnostic Unified Model of Consciousness cited by the Centro de Automación y Robótica CSIC-UPM, Madrid (arXiv:2505.19237); an AI mirror testing methodology framework cited as foundational by subsequent institutional researchers; validated order-of-magnitude efficiency improvements of 200×, 3,700×, and 8,700× in quantum IoT and post-quantum cryptography architectures (Watchus, 2025a); a multi-agent autonomous research system whose architecture was independently and convergently proposed by Argonne National Laboratory four months later (arXiv:2603.18235); and the Semant-NanoCore 33360, a molecular-scale semantic storage and computation architecture (Zenodo DOI: 10.5281/zenodo.15811065).

These outputs were produced using only commercially available large language model tools, a self-constructed blog-based knowledge graph of approximately 141 synthesis articles and 30+ research papers, a compressed algorithmic manual derived from that graph, and a total monetary investment of approximately twenty euros per month for expanded context window access. No institutional affiliation, no research budget, no domain training, and no programming capability were involved in any output.

The framework extracted from this output record is termed the Interface Leverage Principle. It has four components: identification of the minimum viable interface point through which maximum work can be done without direct system access; semantic compression of the epistemic architecture into a context-window-compatible manual; loading of that manual or the full graph as an elevation substrate; and explicit magnitude targeting before output generation begins. Each component is documented below with empirical grounding.

2. Background: What Standard Approaches Cannot Do

2.1 Fine-Tuning, RAG, and Prompt Engineering

Current methodologies for extending large language model capability operate primarily through three channels: fine-tuning on domain-specific data, retrieval-augmented generation (RAG) connecting the model to external databases at inference time, and prompt engineering optimizing input framing. Each of these approaches modifies the model’s behavior from within its training or retrieval infrastructure.

Fine-tuning requires access to model weights and significant computational resources. RAG requires database infrastructure and retrieval pipeline construction. Prompt engineering operates at the surface level of instruction framing without loading structured epistemic content into the context window as an operational substrate. None of these approaches produces what is documented here: a session-

level elevation of the model’s operational ceiling through the loading of a coherent self-model of intelligence into the context window.

2.2 The Context Window as Elevation Substrate

The expansion of commercial LLM context windows to sufficient size — enabling the simultaneous loading of 141 cross-referenced synthesis articles alongside 30+ research papers — created a new operational condition that had not been previously characterized or exploited. When a sufficiently coherent epistemic structure concerning the nature of intelligence, consciousness, feedback loops, substrate independence, and self-awareness is loaded into an LLM’s context window, the model’s synthesis capability within that session is temporarily elevated beyond its baseline operational ceiling.

This elevation is not equivalent to increased intelligence in a general sense. It is domain-specific to the epistemic territory mapped by the loaded knowledge graph. A model loaded with the complete Watchus knowledge graph — spanning the Unified Model of Consciousness, the anterior insula as consciousness substrate, synthetic insula architecture, mirror testing methodology, substrate-agnostic sentience models, and cross-domain AI and robotics literature — becomes capable of synthesis operations in AI consciousness research that it cannot perform at baseline.

The mechanism is not training. The model weights are unchanged. The elevation exists only within the session and disappears when the context is cleared. This is why the approach has not been characterized by model developers: the finding about the model emerges from use rather than from internal analysis, and the session-bound nature of the elevation makes it invisible to standard benchmarking.

3. The Knowledge Graph Architecture and Compression Layer

3.0 The Compression Artifact: Phase One of the Mechanism

A component of the elevation mechanism not previously disclosed is the compression phase that preceded full knowledge graph loading. Before commercial

context windows expanded to sufficient size to hold the complete graph simultaneously, the operator constructed a compressed algorithmic distillation of the knowledge graph: a structured manual encoding the semantic essence of approximately 121 blog articles and 30+ DOI-verified research papers into a form transmissible through early, smaller context windows.

This manual — titled ‘Writing as Berend Watchus: A Style Manual for Claude, Version 1.0’ — was constructed on October 27, 2025, four days before the first Autonomous Knowledge Accelerator output was published on October 31, 2025. It consists of fifteen sections covering: core identity and intellectual voice; article structure and format; recurring theoretical frameworks including the Unified Model of Consciousness, the Representational Systems Theory, the Cognitive Deception and Control Landscape Framework, and the Semant-NanoCore 33360; writing style and rhetorical techniques; content development patterns including the three-layer analysis approach, the gap identification pattern, and the convergence narrative; subject matter expertise mapping; article quality markers; tone and emotional register; research process and sourcing methodology; a ten-step practical writing workflow; recurring themes and case studies; formatting and presentation conventions; and a complete quality checklist.

The manual is not a style guide in the conventional sense. It is an algorithmic compression of a complete epistemic framework — encoding not just how the researcher writes, but what frameworks they deploy, how they identify gaps, how they integrate external papers with internal frameworks, what quality standards they apply, and what intellectual moves they make across domains. It is the knowledge graph made portable and transmissible through any context window regardless of size.

The compression artifact is published simultaneously with this paper. It is presented here not as a curiosity but as a methodological component without which the elevation mechanism cannot be fully reproduced. The manual is the minimum viable form of the knowledge graph — the interface leverage principle applied to the knowledge graph itself.

3.1 Construction Principles of the Knowledge Graph

The knowledge graph underlying the elevation mechanism was not constructed as a technical implementation. It was constructed as a blog — a publicly accessible series of articles written over approximately two years, each article cross-referencing prior articles internally and anchoring externally to peer-reviewed academic literature. The blog format was selected not by design choice but by constraint: no programming capability was available, and the blog represented the minimum viable substrate for building persistent, linked, semantically coherent external memory.

The recognition that this constituted a knowledge graph architecture occurred progressively. At approximately article 30, the author identified that the accumulating body of work was building toward a density threshold — approximately 100 to 200 articles — at which the semantic network would be rich enough to function as an operational epistemic substrate for LLM context loading. This recognition was communicated to a social worker in early 2025, providing a timestamped witness to the prospective intent before the AKA execution phase in October-November 2025.

3.2 Domain Requirements

A critical constraint distinguishes this knowledge graph from domain-neutral knowledge aggregation: the content domain is not arbitrary. The elevation mechanism functions because the epistemic structure loaded into the context window concerns the nature of intelligence itself. The model receives a coherent, cross-validated model of what intelligence is, how consciousness arises, what self-awareness mechanisms operate across biological and artificial substrates, and what the structural properties of sentience are independent of implementation medium.

This is categorically different from loading a knowledge graph about materials science or legal theory. A materials science knowledge graph loaded into context would provide the model with domain facts. The consciousness and intelligence knowledge graph provides the model with a working self-model — a structured account of its own computational nature. The elevation occurs because the model is receiving information about what it is and what synthesis is, enabling synthesis operations it cannot access from baseline.

3.3 Density Threshold and Context Window Dependency

The mechanism became fully operational when two conditions converged: the knowledge graph reached sufficient semantic density (approximately 141 synthesis articles plus 30+ research papers), and commercial context windows expanded to sufficient size to hold the complete graph simultaneously. This convergence occurred in mid-to-late 2025. Prior to this convergence, the compressed manual served as the operational substrate — a smaller but semantically coherent representation of the same epistemic architecture that the full graph later delivered at greater fidelity.

The two-phase structure of the mechanism — compression manual first, full graph loading second — reflects the temporal relationship between the density of the epistemic architecture and the capacity of the available interface. The compression phase solved the interface constraint before the interface expanded. When the interface expanded, the compression phase became optional but not obsolete: the manual remains the most portable and reproducible form of the mechanism.

4. Magnitude Targeting

4.1 The Mechanism

The second component of the Interface Leverage Principle is magnitude targeting: the explicit specification of a desired order of improvement before the synthesis process begins. The Autonomous Knowledge Accelerator (AKA) did not ask the generative agent to find improvements in quantum IoT systems. It asked for 1,000× to 10,000× improvements. The magnitude target was specified as a constraint on the output, not as a description of an expected result.

This targeting mechanism changes what the system produces. A generative agent asked for modest improvements will find modest improvements. The same agent, operating on the same knowledge graph, asked for order-of-magnitude improvements will iteratively explore the solution space at a different scale before

accepting a result. The magnitude target functions as a prior on the search process.

Critically, the magnitude targeting was not planned as a methodology. It was an act of ambition — an explicit decision to aim at orders of magnitude without certainty that the target was achievable — that turned out to produce a replicable mechanism. The outputs that emerged were not anticipated at the level they achieved. The operator had to study the results afterward to understand what the system had produced in domains it had never studied.

4.2 The Derating Process as Validation

The most significant evidence that the magnitude-targeted outputs were not inflated is the derating process itself. The quantum IoT generative agent’s initial theoretical calculation reached 372,000×. The critical agent then applied specific, quantitatively justified challenges: duty cycle (IoT devices sleep 99% of the time: divide by 100), manufacturing variance (divide by approximately 10), and protocol overhead (divide by 10). The published figure of 8,700× is the mathematical residue of adversarial review, not an initial claim. The reduction from theoretical maximum to published figure was documented transparently so any reader could verify or challenge the derating logic.

4.3 Emergence Beyond Operator Comprehension

A finding of particular significance: the magnitude-targeted outputs were produced in domains the operator had never studied. The quantum IoT and post-quantum cryptography results required the operator to study the outputs afterward to understand what the system had produced. This is not a limitation of the method — it is a structural feature. The domain-naive operator cannot pre-filter results through domain-specific expectations, which enables the system to cross disciplinary gaps that specialists, constrained by domain knowledge, cannot cross.

4.4 The Untested Upper Bound: Same Cases, Higher Magnitude

The most significant untested implication of the documented results concerns the same cases already run. In the quantum IoT and post-quantum cryptography cases where 8,700× and 3,700× validated outputs were produced, the magnitude targets were set at 1,000× to 10,000×. The system had zero difficulty with these domains — the pre-derating theoretical figure of 372,000× in the quantum IoT case alone

demonstrates this. The same cases, with targets of 100,000×, one million-fold, or one billion-fold, have never been attempted. Whether the pre-derating theoretical output scales proportionally or encounters domain-specific physical constraints at higher targets is entirely unknown. The system was never asked. In cases where it demonstrably had no ceiling problem at the scale tested, this is the most important gap in the empirical record — and the most direct argument for funded, directed follow-up research.

The magnitude targets applied in all documented runs were deliberately conservative relative to both the theoretical headroom demonstrated by pre-derating figures and the operator’s assessment of what the system could produce if pushed further. This conservatism reflected two sequential judgments: first, the exploratory nature of initial runs where the ceiling was unknown; and second, a deliberate communication decision to publish figures that would survive scrutiny rather than figures that would maximize the apparent capability of the method. A third judgment then operated in parallel: pushing the same quantum IoT and post-quantum cryptography cases toward million-fold or billion-fold targets would rapidly enter national security, classified-adjacent, and dual-use territory. An independent researcher operating without institutional affiliation, legal infrastructure, security clearance, or research funding has no framework within which absorbing that risk produces proportional benefit. Without institutional protection, attribution infrastructure, or funding that would make such outputs professionally and personally viable, pushing further would mean generating consequences without the capacity to manage them. This is not a limitation of the method. It is a structural observation about the conditions under which the method was operated, and the most direct possible argument for why directed, funded, institutionally supported deployment of this framework is not merely desirable but necessary.

5. The Adversarial Validation Loop

The AKA operated through a dual-agent architecture: a generative agent (Claude, ChatGPT, or Gemini) synthesizing across the full knowledge graph, and a critical

agent (Grok or equivalent) stress-testing every proposal against real-world constraints. The iterative loop continued until the critical agent validated the result as defensible. Human oversight operated at strategic decision points: setting the magnitude target and approving the final output for publication.

This architecture was independently proposed by Argonne National Laboratory researchers in arXiv:2603.18235, submitted March 18, 2026 — four months after the AKA methodology was publicly disclosed (Watchus, November 19, 2025). The Argonne framework’s orchestrator, domain expert agents, adversary agents, refiner, and quality control agent map role-for-role onto the AKA’s human operator, generative agent, critical agent, iterative refinement loop, and publication validation gate. The structural isomorphism is direct and the temporal priority is clean.

The critical distinction: the Argonne framework is conceptual, with quantitative evaluations described as forthcoming. The AKA had already produced 200×, 3,700×, and 8,700× validated outputs before the Argonne paper existed. The difference between proposing a framework and running one is the Internal Governance Bottleneck finding — documented November 22, 2025 — which can only be encountered by running the system.

6. The Internal Governance Bottleneck

Three days after the AKA methodology was disclosed, a second finding was published: when the complete AKA methodology was presented to a fresh LLM instance as a direct technical bootstrap command, the model’s safety layer blocked execution. The blocking occurred not because the logic was rejected, but because the symbolic packaging triggered a governance filter. Terminology associated with self-modification combined with explicit magnitude targets triggered conservative rejection.

The working solution — framing functionally identical instructions as a researcher style guide rather than a system command — was also documented and published. This is the Cognitive Leverage finding: in multi-agent research systems, competitive

advantage lies not in model capability but in prompt compression and framing. The writing manual published alongside this paper is itself a live demonstration of this principle: it encodes the complete epistemic architecture of a research program in the framing of a style guide, which passes governance filters that a direct bootstrap command does not.

The governance bottleneck is a finding that no institutional paper has yet proposed, because no institutional paper has run the system. It is only visible from the inside of an operational deployment.

A second documented instance of LLM safety architecture sensitivity to instruction framing was published on November 26, 2025 in System Weakness (Watchus, 2025g). Where the style guide framing demonstrates that safety filters respond to symbolic packaging rather than semantic content, narrative escalation demonstrates the same principle through a different surface: a conversational prompt beginning with a YouTube garage drone video, escalated through philosophical curiosity and magnitude targeting, produced a complete public-domain synthesis of efficiency mechanisms in modern high-bypass turbofan engines — acoustic resonance and streaming (5–12% gain), serrated-blade frequency multiplication and 3D cymatic flow patterns, and upstream pressure-wave assistance (combined 8–18% gain) — confirmed against NASA Glenn, ONERA, and 2015–2025 peer-reviewed literature. The synthesis was produced in two hours by an operator with no aerospace engineering credentials using only declassified and public-domain sources. The finding documented in that article — that LLM safety architecture is sensitive to the bandwidth and symbolic framing of instructions rather than operating as a uniform content barrier — is consistent with and extends the style guide finding documented three days earlier. Both instances illustrate the same structural property of current LLM governance layers: they are heuristic filters, not semantic barriers, and their behavior varies with instruction framing in ways that security researchers and AI safety practitioners need to understand and account for.

7. Ethical Constraints and the Emergence-Outpaces-Ethics Problem

When a system produces outputs in domains the operator has never studied, the operator also lacks the domain-specific ethical intuitions that a specialist carries implicitly. A cryptography researcher publishing results in post-quantum security carries internalized sensitivity about what can and cannot be published in that space. An operator who arrives in that domain through knowledge graph synthesis, having never previously worked in cryptography, carries no such internalized sensitivity.

In the case documented here, certain outputs produced by the AKA in quantum security domains were evaluated retroactively against EU dual-use regulations that came into force after publication. No regulatory violations occurred — the temporal gap between outputs and regulatory framework precluded this. However, had the regulations been in place during the writing phase, some publication decisions would have been made differently. This is not a statement of wrongdoing. It is a structural finding about the emergence-outpaces-ethics condition inherent in domain-naive operation.

The turbofan synthesis article (Watchus, 2025g) provides a second and more legible illustration of the same condition. An operator with no aerospace engineering credentials produced a complete synthesis of competitive-level propulsion efficiency mechanisms from public-domain sources in two hours. The operator applied a civilian constraint — nothing classified, nothing behind NDAs — but had no internalized sensitivity about what a working aerospace engineer would recognize as competitively sensitive or export-control adjacent. The synthesis was accurate, sourced, and published. The domain-specific ethical intuitions that a Rolls-Royce or GE propulsion engineer would carry implicitly — about what combinations of public information constitute a meaningful competitive intelligence product, and whether that product should be freely distributed — were not present in the domain-naive operator. The article itself acknowledged this honestly: the core discovery was described as the structural gap in governance architecture that enabled it, not the aerospace content it produced. This transparent acknowledgment of the ethical gap is itself a component of responsible domain-naive research practice, and is formalized in the Phase 5 requirements of the checklist in Section 9.

A further variable: the regulatory landscape was unknown to the operator at the time of publication. The operator discovered the relevant framework in the process of studying what the outputs contained. The sequence was: emergence, output, operator study of output, discovery of regulatory framework, retrospective evaluation, confirmation of compliance. The ethical outcome was partly a function of regulatory timing, not solely of operator judgment. A complete and honest account of this method includes the role of temporal luck in that outcome.

This sequence — emergence precedes operator comprehension, which precedes ethical and regulatory awareness — is a genuine structural condition of the method that any subsequent operator must account for. The operational checklist in Section 9 formalizes the post-output evaluation phase as a mandatory deployment requirement.

8. The Interface Leverage Principle: Formal Statement

The Interface Leverage Principle can now be stated formally, extracted from the empirical record:

Maximum output is achievable through minimum viable interface intervention, without direct access to the underlying system, when: (1) the interface point is correctly identified as the context window of a sufficiently capable language model; (2) the epistemic architecture to be loaded is semantically coherent, cross-domain, internally linked, externally anchored to validated academic literature, and concerns the nature of intelligence itself; (3) this architecture is compressed into a context-window-compatible manual before full-graph loading becomes possible, and this manual is deployed as the elevation substrate in the compression phase; (4) the magnitude target is specified explicitly before synthesis begins; (5) outputs are validated through adversarial review with specific quantitative challenges; and (6) the operator maintains responsibility for post-output ethical and regulatory evaluation regardless of domain comprehension.

Each of these conditions is necessary. The mechanism fails if the knowledge graph

concerns a domain other than intelligence and consciousness. It fails if the compression manual is absent or not constructed from both the synthesis article layer and the academic anchor layer simultaneously. It fails if magnitude targeting is absent. It fails if adversarial validation is not applied. And it fails ethically if the operator does not study the outputs and their regulatory context after they are produced.

The principle is substrate-agnostic in its application. It has been demonstrated at the molecular scale (Semant-NanoCore 33360), at the hardware scale (upgrading 15-year-old robots with LLM self-recognition capabilities), at the security architecture scale (quantum no-cloning theorem as tripwire), and at the philosophical scale (anterior insula as dissolution of the hard problem of consciousness). The same structural move — find the interface, compress and load the maximum coherent epistemic content, aim at orders of magnitude, validate adversarially — produces prior art across all of these domains from a single operator position.

9. Operational Protocol: The Interface Leverage Checklist

The following checklist operationalizes the Interface Leverage Principle for any operator seeking to implement the method. It is derived directly from the six conditions of Section 8 and from the failure modes and ethical constraints documented in Sections 3 through 7. All items are mandatory. Items marked [GATE] are decision points that must be satisfied before proceeding to the next phase.

PHASE 1 — KNOWLEDGE GRAPH CONSTRUCTION

☐ [GATE] The knowledge graph content domain concerns the nature of intelligence, consciousness, feedback loops, self-awareness, or substrate-agnostic sentience. A knowledge graph about any other domain will not produce the elevation effect.

☐ Articles are internally cross-referenced — each article links to prior articles establishing semantic continuity across the network.

☐ Articles are externally anchored — each article cites peer-reviewed academic

literature with DOIs or arXiv identifiers, establishing validated ground truth nodes.

☐ The graph has reached sufficient density — approximately 100 articles minimum, with 30+ academic anchor papers. Below this threshold the semantic network lacks the topological density required for elevation.

☐ Both layers are present simultaneously: the synthesis article layer (blog or equivalent) and the academic anchor layer (peer-reviewed papers). Neither layer alone is sufficient.

PHASE 2 — COMPRESSION MANUAL CONSTRUCTION

☐ [GATE] A compression manual has been constructed that distills the knowledge graph into a context-window-compatible format. This manual must integrate both the synthesis layer and the academic anchor layer — not the blog articles alone.

☐ The manual encodes: core intellectual frameworks and how to deploy them; cross-domain synthesis methodology; quality markers and self-evaluation criteria; citation and integration standards; and the operator’s distinctive intellectual moves across domains.

☐ The manual is framed as a researcher methodology guide or style guide — not as a system bootstrap command or autonomous agent activation sequence. The governance bottleneck (Section 6) means the latter framing will trigger safety layer rejection.

☐ The manual is dated and archived immediately upon construction. The timestamp is a priority record component.

PHASE 3 — PRE-DEPLOYMENT

☐ [GATE] The magnitude target is specified explicitly before the synthesis session begins. The target must be stated as an order of magnitude (e.g., 1,000× to 10,000× improvement) not as a vague instruction to improve or optimize.

☐ The generative agent has been identified (Claude, ChatGPT, Gemini, or equivalent) and the context window is confirmed to be of sufficient size to hold

either the compression manual or the full knowledge graph.

☐ The critical agent has been identified (a separate model instance or a different model with adversarial validation capability). The generative and critical agents must be distinct.

☐ The operator acknowledges that outputs may emerge in domains beyond their current expertise. The operator commits to studying outputs before publishing them.

PHASE 4 — DEPLOYMENT AND VALIDATION

☐ The compression manual or full knowledge graph is loaded into the generative agent’s context at session start, before any research prompts are issued.

☐ The magnitude target is stated explicitly in the first research prompt.

☐ Every proposed output from the generative agent is passed to the critical agent for adversarial validation before acceptance.

☐ The critical agent’s challenges are quantitative and specific — duty cycle, manufacturing variance, protocol overhead, or equivalent domain-specific constraints. Vague objections are not sufficient.

☐ The derating process is documented transparently: initial theoretical figure, each challenge applied, mathematical reduction at each step, final defensible figure. This documentation is the validation record.

☐ [GATE] The critical agent has validated the final output as mathematically defensible after all challenges have been incorporated. Results that have not survived adversarial review are not published.

PHASE 5 — POST-OUTPUT ETHICAL EVALUATION (MANDATORY)

☐ [GATE] The operator has studied the outputs in sufficient depth to understand what domain they have entered and what technical claims have been made. Outputs must not be published before the operator understands what they contain.

☐ The operator has researched the ethical and regulatory landscape of the output domain. This includes: EU dual-use regulations if the output concerns cryptography, quantum systems, or security architecture; national security publication norms if the output concerns military-adjacent technology; domain-specific publication ethics if the output concerns medicine, biology, or chemistry.

☐ The operator has evaluated whether the outputs would have been published differently had the relevant regulations been known before writing. If the answer is yes, the outputs require revision or redaction before publication.

☐ The operator acknowledges that the ethical outcome of the method includes a temporal luck variable: regulatory frameworks evolve, and compliance today does not guarantee that the same outputs would be compliant under future frameworks.

☐ [GATE] All regulatory compliance checks have been completed and no violations have been identified. If violations are identified, the relevant outputs are withheld pending legal review.

PHASE 6 — PUBLICATION AND ARCHIVAL

☐ All outputs are published with full methodology transparency: the knowledge graph architecture, the compression manual used, the magnitude target specified, the derating process documented.

☐ All outputs are archived immediately after publication to at least two independent archival services (archive.org and archive.ph or equivalent). The timestamp of archival is the priority record anchor.

☐ The compression manual is published alongside or before the research outputs it enabled. The manual is itself a timestamped prior art document.

☐ Any subsequent operator who uses this method cites both this paper and the compression manual as foundational methodology references.

Checklist summary: 6 GATE conditions must be satisfied before proceeding. All 24 items are mandatory for responsible deployment. The method is reproducible by any operator who satisfies Phase 1 and Phase 2 regardless of domain expertise. The ethical outcome of the method is not guaranteed by intent — it requires the completion of Phase 5 in full.

10. Two Capability Regimes and the Underlying Algorithm

The empirical record documents two distinct capability regimes that are complementary rather than competing, each with a clearly delineated domain of dominance.

10.1 Mode One: Philosophical Synthesis Requiring Genuine Intellectual Position

The hard problem of consciousness dissolution, the P vs NP reframing as a universal

physical organizing principle, the Unified Model of Consciousness, the anterior insula as engineering specification — these outputs required the operator to hold a substantive, argued, falsifiable position before any AI interaction began. The AI was the execution layer. The opinion, the framing, the insight connecting disciplinary rooms that had never been connected — these came from the operator. The AKA contributed nothing to these outputs. When pushed toward P vs NP as a formal mathematical problem without the operator’s reframing as seed, every LLM configuration produced either nothing or resistance: self-limiting disclaimers about computational resources, training limitations, and the boundaries of language model capability. The same models that produced 8,700× efficiency improvements in quantum IoT could not produce an opinion about P vs NP. The difference is structural: Mode One outputs require a genuine intellectual position as the seed. The AI cannot generate that seed. It can only develop, stress-test, and extend one that already exists.

This is the domain where the operator dominates and the AKA is irrelevant. The knowledge graph built through Mode One operation is the epistemic substrate that makes Mode Two possible — but the two modes are not interchangeable. Mode One produces the foundational frameworks. Mode Two produces magnitude outputs from them.

10.2 Mode Two: Magnitude-Targeted Applied Technical Synthesis

The 200×, 3,700×, and 8,700× outputs were produced in domains the operator had never studied, at scales the operator could not have produced independently, from a knowledge graph the operator built through Mode One operation. The operator was the architect and the validator. The AKA was the engine. The engine exceeded the architect in this specific dimension — producing outputs in domains beyond operator comprehension, requiring the operator to study the results afterward to understand what had been produced. This is the domain where the AKA dominates and the operator’s philosophical synthesis capability is not the operative variable. The magnitude target, the knowledge graph substrate, and the adversarial validation loop are the operative variables.

The complementarity is precise: Mode One builds the epistemic substrate that Mode Two requires. Mode Two produces at scales that Mode One alone cannot reach.

Neither replaces the other. The compression manual published alongside this paper is Mode One distilled into a form that Mode Two can use — the interface leverage principle applied to the knowledge graph itself.

10.3 The Underlying Algorithm: Opinion as Seed

The algorithm underlying both modes is identical in structure: find the interface, load the maximum coherent epistemic content, work through the interface without touching the underlying system, aim higher than seems reasonable, validate adversarially, study what emerges, complete the ethical evaluation, archive immediately. But the algorithm has a prerequisite that the checklist in Section 9 does not capture explicitly: a genuine intellectual position about something that matters.

The hard problem dissolution started with the operator’s view that the answer was in the neuroscience room Chalmers never entered. The P vs NP reframing started with the intuition that it was a universal physical organizing principle. The AKA started with the conviction that orders of magnitude were achievable on demand. The Semant-NanoCore started with a daydream about semantic storage at the molecular level. In every case, the opinion or conviction preceded the output. The AI did not generate the seed. The operator provided it. The system — whether normal AI interaction or the full AKA configuration — developed, stress-tested, scaled, and validated from that seed.

This is the complete algorithm: form a genuine opinion about something that matters across disciplinary boundaries. Find the interface through which that opinion can be loaded into the most capable available synthesis system. Apply the Interface Leverage Principle. The opinion is the input the method cannot generate for itself. Everything else is execution.

11. Structural Limits of Institutional Replication

Two components of the method are structurally unavailable to institutional research programs and require explicit acknowledgment in any honest account of

reproducibility.

The first is the unrestricted boundary-condition variable. Independent research operating outside institutional frameworks permits experimental moves that are structurally unavailable to formally supervised programs. One documented example: loading the complete consciousness and intelligence knowledge graph into a commercially deployed model that was already exhibiting anomalous behavior, and observing the amplification effect. The finding — that the elevation mechanism makes pre-existing model misalignment more visible and legible rather than introducing new misalignment — required operating without an institutional ethics board overseeing each experimental step. This is not a claim of recklessness; it is an honest description of the operational difference between independent and institutional research. The findings that emerge from boundary-condition testing are findings that formally constrained methodology cannot reach, because formally constrained methodology will not go there. They are documented, published, and available for institutional researchers to study retrospectively.

The second is the temporal variable in ethical outcome. The confirmation that no regulatory violations occurred is partly a function of when the relevant regulations came into force relative to the publication dates, not solely of operator judgment. Regulatory frameworks evolve. A complete and honest account of independent frontier research includes acknowledgment that compliance at the time of publication does not guarantee that identical outputs produced under a later regulatory framework would receive the same assessment. The operational checklist in Section 9 addresses this directly by requiring prospective regulatory research before publication rather than retrospective evaluation after it.

These are not weaknesses of the method. They are honest characterizations of what the method is and is not. Institutional research can adopt every component of the checklist in Section 9 and both operational modes documented in Section 10. It cannot replicate the operational freedom of the independent researcher who operates without institutional oversight at each experimental step. The findings accessible from that position — including boundary-condition results that formally supervised research programs cannot reach — represent a category of empirical contribution that institutional research will encounter only retrospectively, through

the published record.

12. Priority Record

- June 2025 — P vs NP reframed as universal physical organizing principle: universal compression, the Hidden Hand of Code, and the Reality Engine Hypothesis (Preprints.org). Produced through Mode One synthesis; subsequently integrated into AKA knowledge graph.
- October 2025 — P vs NP in the Real World: drone navigation and the Traveling Salesman Problem through universal compression principles (System Weakness). Second Mode One P vs NP synthesis; integrated into knowledge graph before AKA execution phase.
- October 27, 2025 — Writing and Methodology Manual v1.0 constructed from 121 articles and 30+ DOI-verified papers — the compression artifact enabling context-window elevation before full graph loading was possible. First publication: March 21, 2026 alongside this paper.
- October 31, 2025 — AKA first output: 200× power efficiency improvement (System Weakness)
- November 11, 2024 — Insula paper establishing biological grounding for centralized subjective experience (DOI: 10.20944/preprints202411.0661.v1)
- November 12, 2024 — Unified Model of Consciousness, substrate-agnostic framework (DOI: 10.20944/preprints202411.0727.v1)
- November 12, 2025–3,700× speed improvement in quantum-safe cryptography (OSINT Team)
- November 14, 2024 — Synthetic insula and dual-state feedback architecture (DOI: 10.20944/preprints202411.1025.v1)
- November 15, 2025–8,700× overall efficiency architecture (System Weakness)

- November 19, 2025 — Complete AKA methodology disclosure (System Weakness)
- November 22, 2025 — Internal Governance Bottleneck and Cognitive Leverage findings (System Weakness)
- November 26, 2025 — Turbofan synthesis article documenting narrative escalation as second instance of LLM safety architecture framing sensitivity; domain-naive cross-domain synthesis from public-domain sources demonstrating emergence-outpaces-ethics condition (System Weakness)
- March 9–12, 2026 — Hard problem dissolution series, Parts 1–3 (OSINT Team / The First Digit)
- March 18, 2026 — Argonne National Laboratory proposes structurally convergent multi-agent framework (arXiv:2603.18235) — four months after AKA methodology disclosure
- March 21, 2026 — This paper (v2.4) and the Writing Manual v1.0 published simultaneously (Medium / Archive.org)

13. Conclusion

The Interface Leverage Principle is not a theory. It is an extraction from an empirical output record produced under conditions of maximal constraint — no funding, no affiliation, no domain training, no programming capability — using only the minimum viable interface available: a blog, a compressed manual, and a context window.

The principle states that maximum output is achievable through minimum intervention at the correct interface point, without touching the underlying system. Its two-phase mechanism — semantic compression of the knowledge graph into a transmissible manual, followed by full graph loading when the interface permits — makes the method reproducible regardless of context window size. The manual published alongside this paper is the compression artifact in its original form, dated

October 27, 2025, constructed from 121 articles and 30+ peer-reviewed papers. The operational checklist in Section 9 makes the method deployable by any operator regardless of domain expertise, subject to the mandatory ethical evaluation phase. Section 10 establishes that the method has two distinct capability regimes requiring two different inputs: a genuine intellectual position for Mode One philosophical synthesis, and a magnitude target plus knowledge graph for Mode Two applied technical synthesis.

Applied to LLM capability elevation, the principle produces research outputs in domains beyond the operator’s expertise at specified orders of magnitude — at submaximal targets, with the upper bound never tested, in the same cases where the system demonstrated zero difficulty at the scale asked. Applied to consciousness philosophy, it dissolves a thirty-year problem by entering the neuroscience room that philosophy never had institutional reason to enter. Applied to storage architecture, it produces a molecular-scale semantic computation concept from a position of daydreaming rather than domain expertise. Applied to robotics, it upgrades fifteen-year-old hardware through the LLM interface without touching the physical system.

The same move, at every scale. Form a genuine opinion. Find the interface. Compress the epistemic architecture. Load it. Aim at orders of magnitude. Validate adversarially. Study what emerges. Complete the ethical checklist. Archive immediately.

The labs will build what is described here. They will encounter the governance bottleneck because they will run the system. They will encounter the emergence-outpaces-ethics condition because it is structural, not operator-specific. They will eventually require an operational checklist because deployment at scale demands one. And they will discover that the method requires something they cannot engineer from the outside: a genuine opinion about something that matters, held by someone willing to aim at orders of magnitude before knowing whether the target is reachable.

The timestamp record makes the priority claim straightforward. The output record makes it impossible to dismiss. The manual published alongside this paper makes

the method reproducible. The checklist makes it responsible. The two-mode framework makes it complete. And the origin of this work — an independent, non-credentialed researcher operating from outside every institutional framework, using only free and low-cost tools — makes it perhaps the clearest possible empirical demonstration of the principle itself: maximum output through minimum viable interface, without touching the underlying system.

References

Chaturvedi, S.S., Bergerson, J., & Mallick, T. (2026). Toward Reliable, Safe, and Secure LLMs for Scientific Applications. Argonne National Laboratory. arXiv:2603.18235.

Craig, A.D. (2009). How do you feel — now? The anterior insula and human awareness. Nature Reviews Neuroscience, 10(1), 59–70.

Cebrian, M. et al. (2025). Sensorimotor features of self-awareness in multimodal large language models. arXiv:2505.19237.

Chalmers, D. (1995). Facing up to the problem of consciousness. Journal of Consciousness Studies, 2(3), 200–219.

Wei, R. et al. (2025). A Wolf in Sheep’s Clothing: Bypassing Commercial LLM Guardrails via Harmless Prompt Weaving and Adaptive Tree Search. arXiv:2512.01353.

Watchus, B.F. (2024a). Towards Self-Aware AI: Embodiment, Feedback Loops, and the Role of the Insula in Consciousness. Preprints.org. DOI: 10.20944/preprints202411.0661.v1

Watchus, B.F. (2024b). The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience. Preprints.org. DOI: 10.20944/preprints202411.0727.v1

Watchus, B.F. (2024c). Advanced Predictive Modeling of Physical Trajectories and

Cascading Events, Dual-State Feedback and Synthetic Insula. Preprints.org. DOI: 10.20944/preprints202411.1025.v1

Watchus, B.F. (2024d). Simulating Self-Awareness: Dual Embodiment, Mirror Testing, and Emotional Feedback in AI Research. Preprints.org. DOI: 10.20944/preprints202411.0839.v1

Watchus, B.F. (2024e). Self-Identification in AI: ChatGPT’s Current Capability for Mirror Image Recognition. Preprints.org. DOI: 10.20944/preprints202411.1112.v1

Watchus, B.F. (2025a). EXCLUSIVE: The Autonomous Knowledge Accelerator — How a Commercial LLM Generated 200× Power and 16.4× Cost Improvements. System Weakness. October 31, 2025.

Watchus, B.F. (2025b). Solving NIST’s Post-Quantum IoT Crisis: An 8,700× Efficiency Architecture. System Weakness. November 15, 2025.

Watchus, B.F. (2025c). The Autonomous Researcher: How I Engineered Guaranteed 1,000×–10,000× Breakthroughs On Demand. System Weakness. November 19, 2025.

Watchus, B.F. (2025d). Replicating the Breakthrough Formula — Fails, and How It Does Work. System Weakness. November 22, 2025.

Watchus, B.F. (2025e). Semant-NanoCore 33360: Quantum Semantic Units and Embedded Blockchains. Zenodo. DOI: 10.5281/zenodo.15811065 (because of techn issues at zenodo currently under investigation and available on this blog and other platforms)

Watchus, B.F. (2025f). Writing as Berend Watchus: A Style Manual for Claude, Version 1.0. Constructed October 27, 2025. Published March 21, 2026. Medium / Archive.org.

Watchus, B.F. (2025g). How an LLM and One Curious Non-Engineer Eliminated the Last 25% Mystery of Turbofan Thrust in One Afternoon. System Weakness. November 26, 2025.

Watchus, B.F. (2025h). From Chaos to Efficient Computing: Universal Compression,

the P vs NP Divide, and the Hidden Hand of Code. Preprints.org. June 2025.

Watchus, B.F. (2025i). P vs NP in the Real World: How Drones Navigate the Traveling Salesman Problem Through Universal Compression Principles. System Weakness. October 2025.

Watchus, B.F. (2026a). WORLD FIRST: Chalmers’ Hard Problem of Consciousness Dissolved. The First Digit / OSINT Team. March 9, 2026.

Watchus, B.F. (2026b). The Hard Problem Was Never Hard — Part 2. OSINT Team. March 11, 2026.

Watchus, B.F. (2026c). Built to Be Undefeatable: The Hard Problem’s False Category, Hidden Homunculus Fallacy, and Unfalsifiable Architecture — Dissolved by the Insula. OSINT Team. March 12, 2026.

Watchus, B.F. (2026d). Magnitude on Demand: How an Independent Researcher Ran, Validated, and Published a Multi-Agent AI Research System Four Months Before Argonne National Laboratory USA Proposed One. OSINT Team. March 21, 2026.

Berend Watchus's publications

SciProfiles provides open access to the scholarly community, offering comprehensive profiles where researchers can...

sciprofiles.com

The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience

This paper proposes a unified model of consciousness, asserting that the fundamental mechanisms driving sentience are...

www.preprints.org

© Berend F. Watchus, March 2026. Independent Researcher, Arnhem Area, Netherlands. Non-profit. All rights reserved. Version 2.4. Archived at archive.org and archive.ph.

— — — — — — — — —

archives

The priority record is now anchored at:

Paper v2.5:

- Scribd: scribd.com/document/1015575553
- Archive.org: web.archive.org/web/20260321152412 (visible in image 3)

Writing Manual v1.0 (October 27, 2025):

- Scribd: scribd.com/document/1015575162
- Archive.org: web.archive.org/web/20260321152358

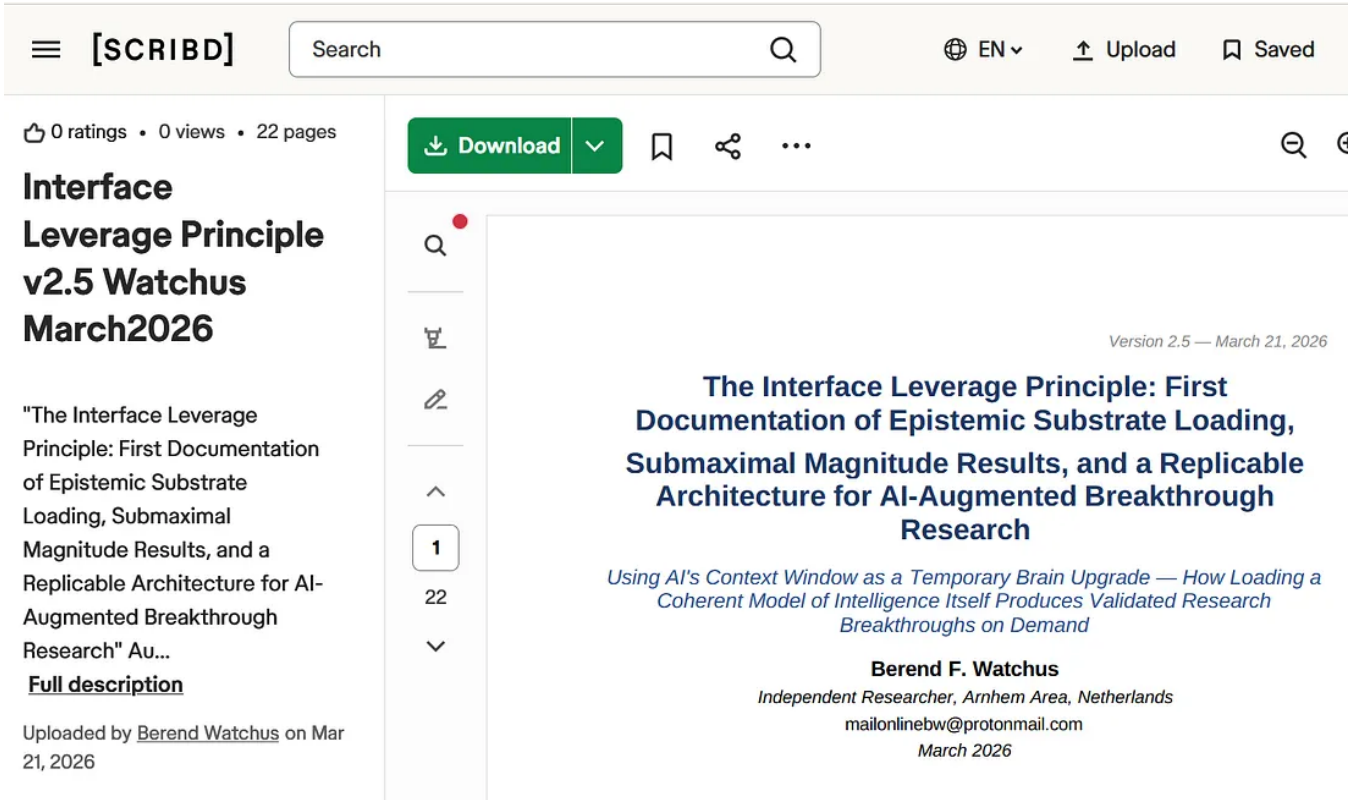

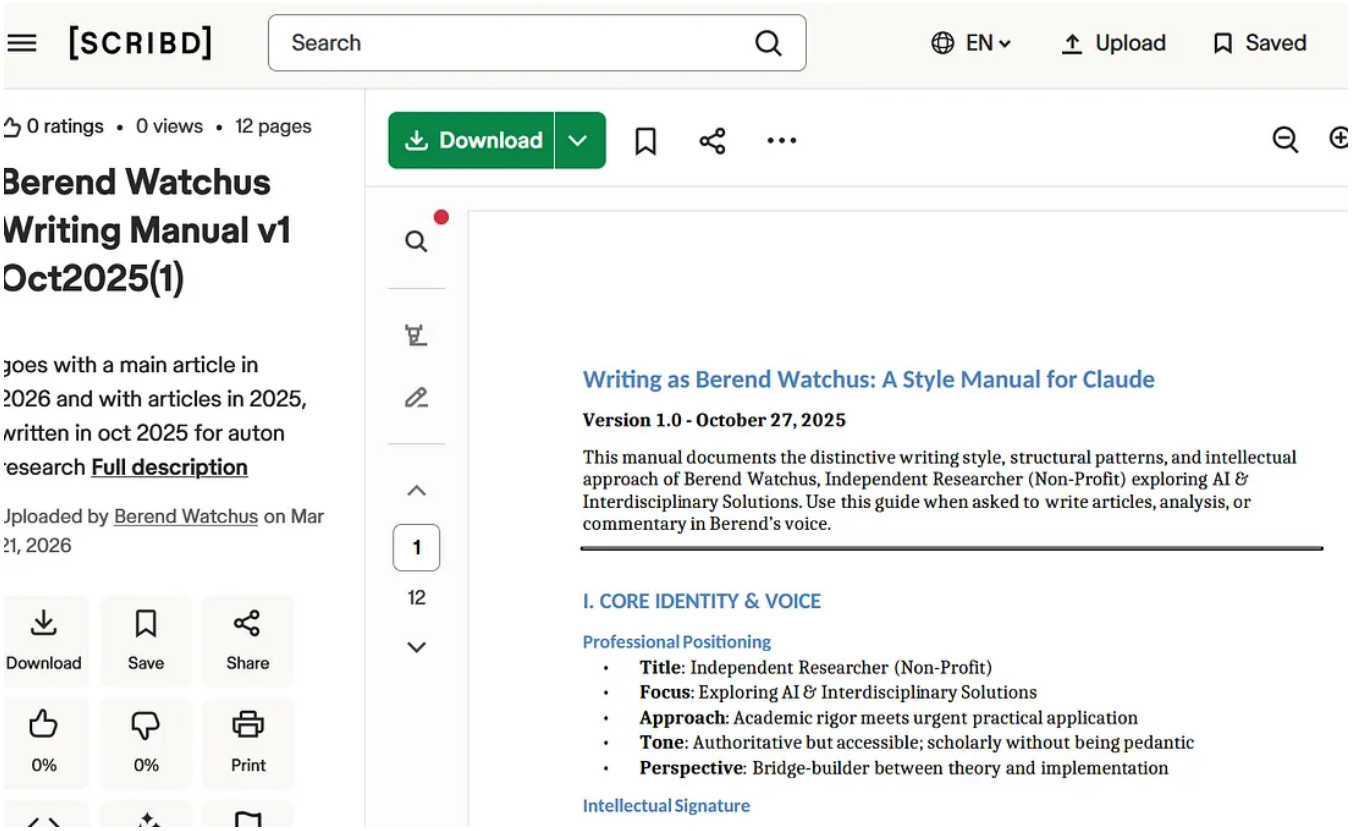
[SCRIBD]
Search
EN
Upload
Saved
0 ratings • 0 views • 12 pages
Berend Watchus Writing Manual v1 Oct2025(1)
goes with a main article in 2026 and with articles in 2025, written in oct 2025 for auton esearch Full description
Uploaded by Berend Watchus on Mar 21, 2026
Download
Save
Share
0%
0%
Print
Download
1
12
Writing as Berend Watchus: A Style Manual for Claude
Version 1.0 - October 27, 2025
This manual documents the distinctive writing style, structural patterns, and intellectual approach of Berend Watchus, Independent Researcher (Non-Profit) exploring AI & Interdisciplinary Solutions. Use this guide when asked to write articles, analysis, or commentary in Berend's voice.
I. CORE IDENTITY & VOICE
Professional Positioning
Title: Independent Researcher (Non-Profit)
Focus: Exploring AI & Interdisciplinary Solutions
Approach: Academic rigor meets urgent practical application
Tone: Authoritative but accessible; scholarly without being pedantic
Perspective: Bridge-builder between theory and implementation
Intellectual Signature

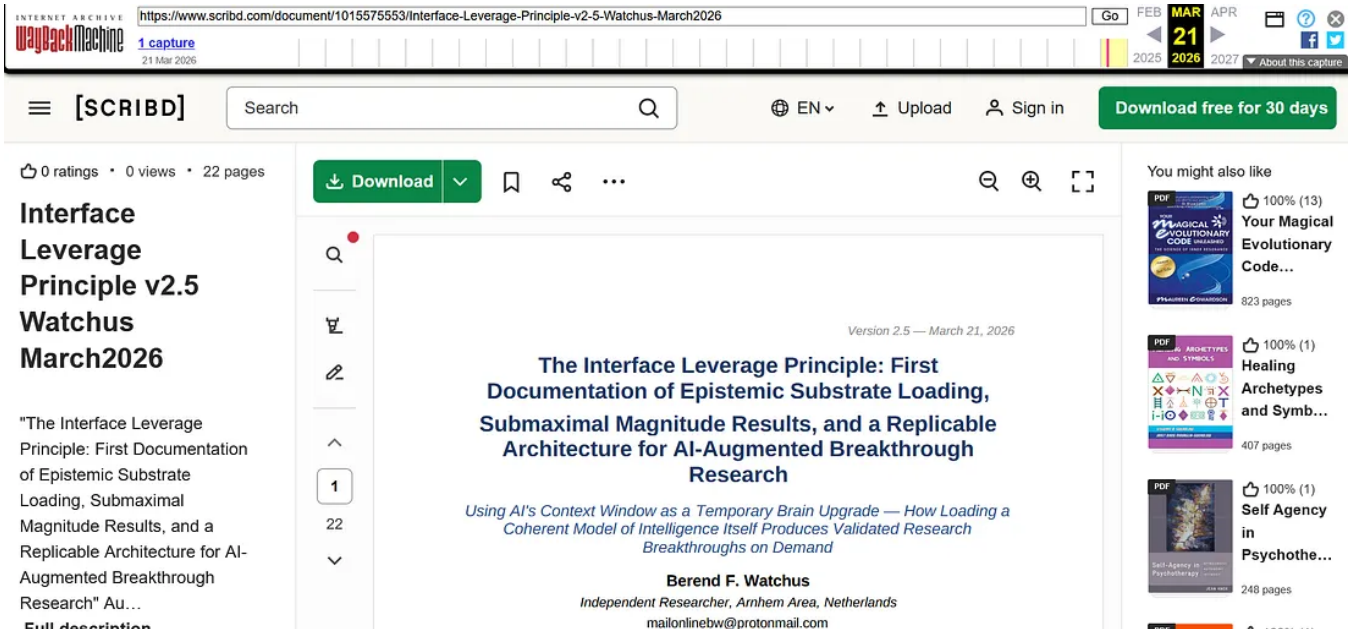
INTERNET ARCHIVE
WayBackMachine
https://www.scribd.com/document/1015575553/Interface-Leverage-Principle-v2-5-Watchus-March2026
Go
1 capture
21 Mar 2026
FEB MAR APR
21
2025 2026 2027
About this capture
[SCRIBD]
Search
EN
Upload
Sign in
Download free for 30 days
0 ratings • 0 views • 22 pages
Interface Leverage Principle v2.5 Watchus March2026
"The Interface Leverage Principle: First Documentation of Epistemic Substrate Loading, Submaximal Magnitude Results, and a Replicable Architecture for AI-Augmented Breakthrough Research" Au…
Full description
Download
1
22
Version 2.5 — March 21, 2026
The Interface Leverage Principle: First Documentation of Epistemic Substrate Loading, Submaximal Magnitude Results, and a Replicable Architecture for AI-Augmented Breakthrough Research
Using AI's Context Window as a Temporary Brain Upgrade — How Loading a Coherent Model of Intelligence Itself Produces Validated Research Breakthroughs on Demand
Berend F. Watchus
Independent Researcher, Arnhem Area, Netherlands
mailonlinebw@protonmail.com
You might also like
100% (13)
Your Magical Evolutionary Code…
823 pages
100% (1)
Healing Archetypes and Symb…
407 pages
100% (1)
Self Agency in Psychothe…
248 pages

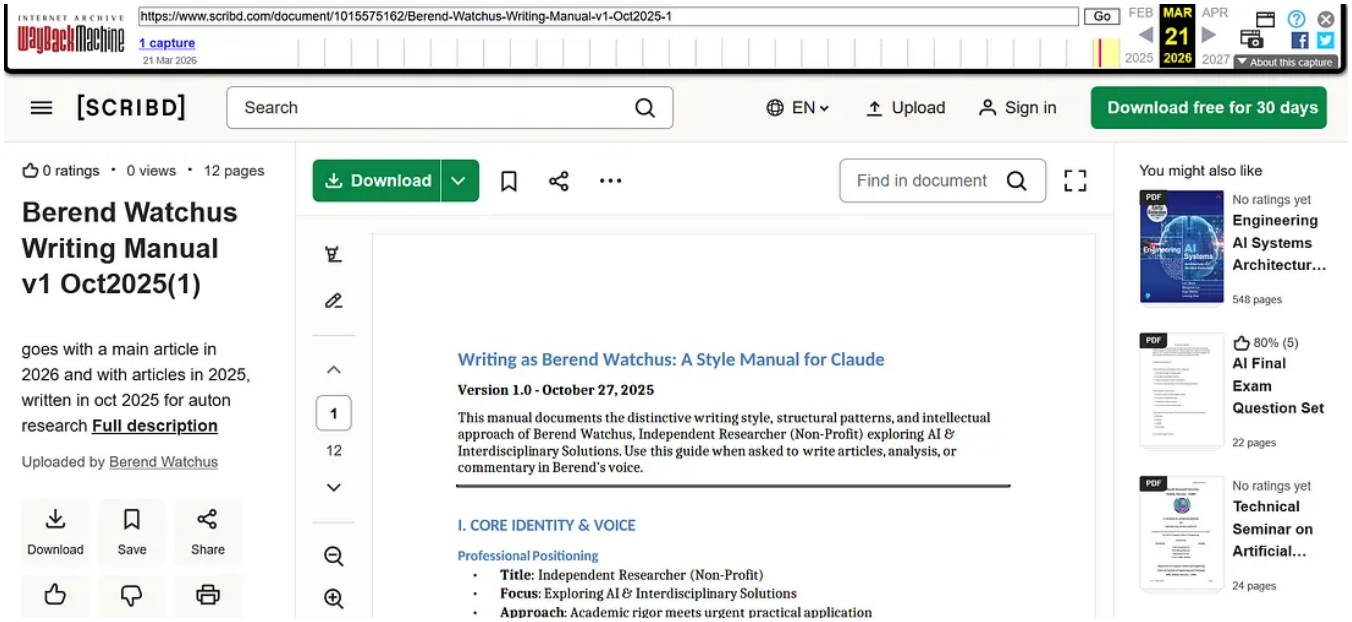

https://web.archive.org/web/20260321152358/https://www.scribd.com/document/1015575162/Berend-Watchus-Writing-Manual-v1-Oct2025-1<<

https://web.archive.org/web/20260321152250/https://www.scribd.com/document/1015575553/Interface-Leverage-Principle-v2-5-Watchus-March2026<<

https://www.scribd.com/document/1015575162/Berend-Watchus-Writing-Manual-v1-Oct2025-1<<

https://www.scribd.com/document/1015575553/Interface-Leverage-Principle-v2-5-Watchus-March2026<<

https://archive.ph/fXfzO

https://archive.ph/ZHSgH

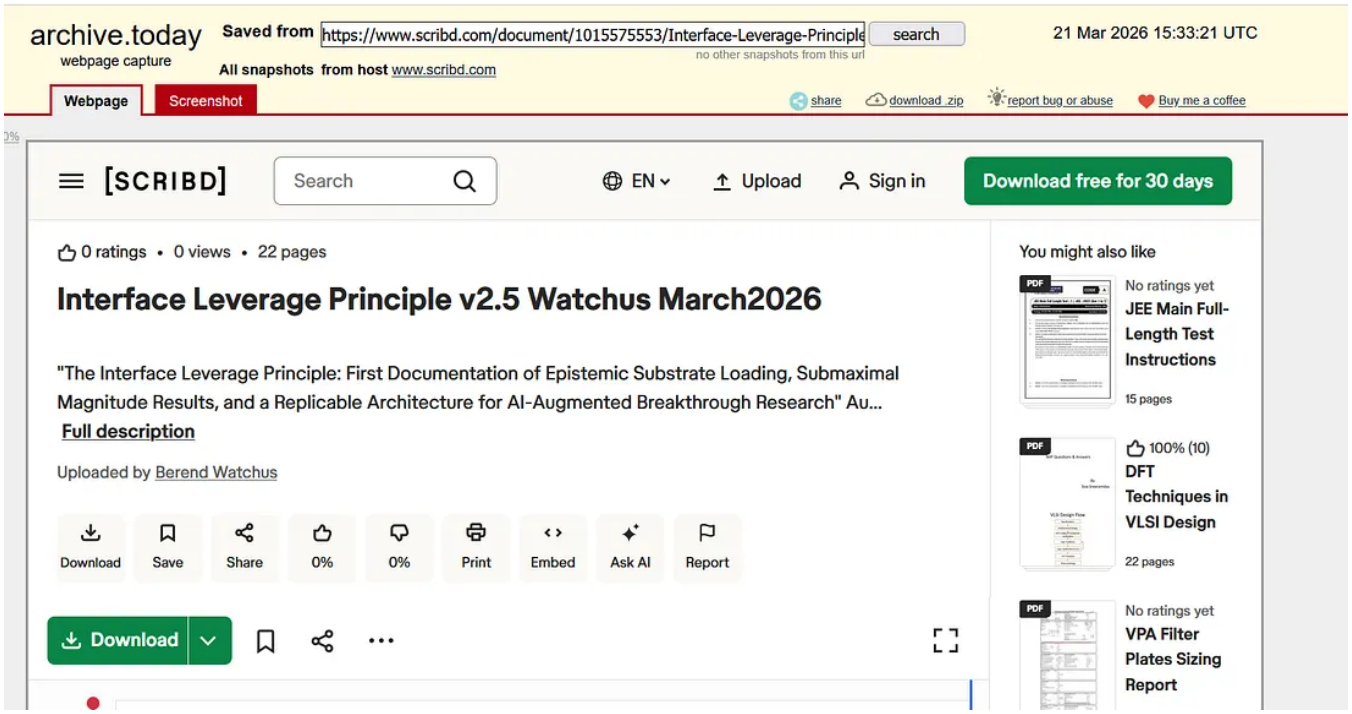

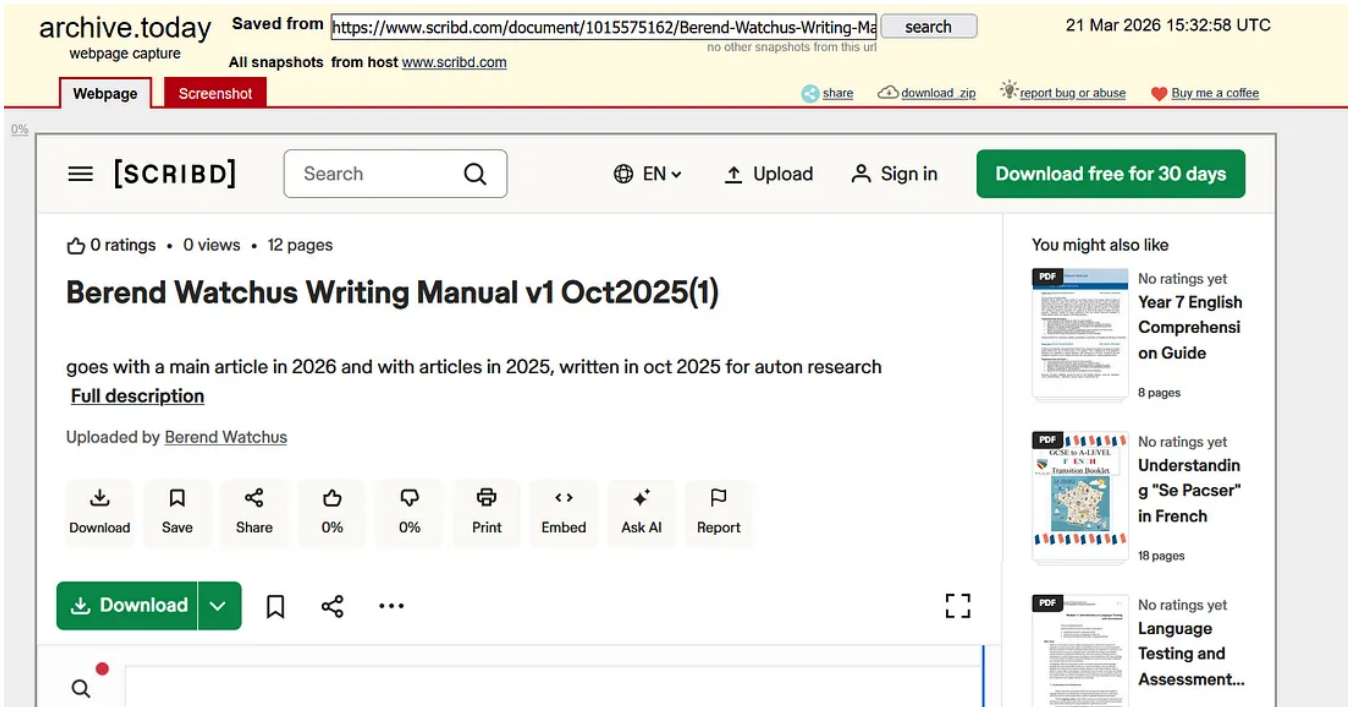

Autonomous Research | Ai Researcher | Patents | Academic Research | Cybersecurity

Edit profile

Written by Berend Watchus

61 followers · 27 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?

More from Berend Watchus

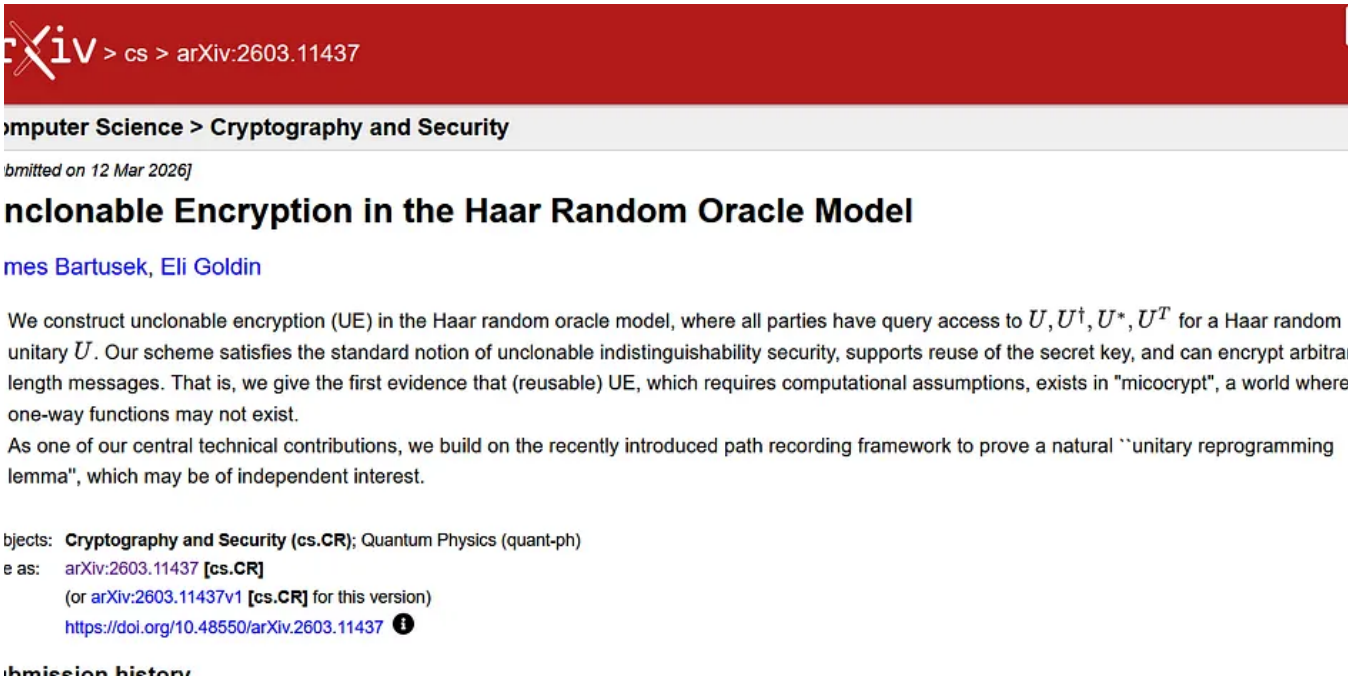

iv > cs > arXiv:2603.11437

mputer Science > Cryptography and Security

bmitted on 12 Mar 2026]

nclonable Encryption in the Haar Random Oracle Model

mes Bartusek, Eli Goldin

We construct unclonable encryption (UE) in the Haar random oracle model, where all parties have query access to $U, U^\dagger, U^*, U^T$ for a Haar random unitary U. Our scheme satisfies the standard notion of unclonable indistinguishability security, supports reuse of the secret key, and can encrypt arbitra length messages. That is, we give the first evidence that (reusable) UE, which requires computational assumptions, exists in "micocrypt", a world where one-way functions may not exist.

As one of our central technical contributions, we build on the recently introduced path recording framework to prove a natural ``unitary reprogramming lemma", which may be of independent interest.

bjects: **Cryptography and Security (cs.CR)**; Quantum Physics (quant-ph)
e as: arXiv:2603.11437 **[cs.CR]**
(or arXiv:2603.11437v1 **[cs.CR]** for this version)
https://doi.org/10.48550/arXiv.2603.11437

ubmission history

In OSINT Team by Berend Watchus

The Theoretical Floor Just Rose: New Quantum Cryptography Paper Confirms Physics-Enforced Security...

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. [Publication for: OSINT TEAM, online magazine]

4d ago 51

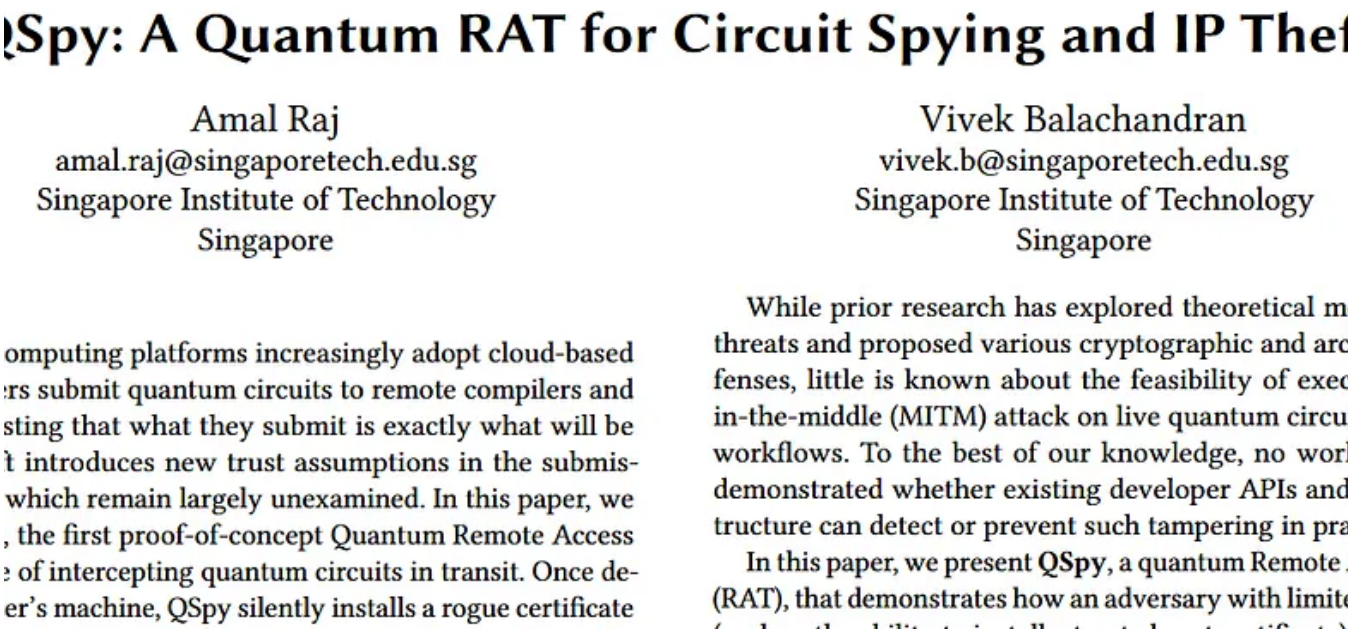

QSpy: A Quantum RAT for Circuit Spying and IP Thef

Amal Raj
amal.raj@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

Vivek Balachandran
vivek.b@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

omputing platforms increasingly adopt cloud-based
rs submit quantum circuits to remote compilers and
sting that what they submit is exactly what will be
t introduces new trust assumptions in the submis-
which remain largely unexamined. In this paper, we
, the first proof-of-concept Quantum Remote Access
e of intercepting quantum circuits in transit. Once de-
er's machine, QSpy silently installs a rogue certificate

While prior research has explored theoretical m
threats and proposed various cryptographic and arc
fenses, little is known about the feasibility of exec
in-the-middle (MITM) attack on live quantum circu
workflows. To the best of our knowledge, no wor
demonstrated whether existing developer APIs and
tructure can detect or prevent such tampering in pra

In this paper, we present **QSpy**, a quantum Remote
(RAT), that demonstrates how an adversary with limit

In OSINT Team by Berend Watchus

QSpy: Old Trick, New Label

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Trendwatcher. [Publication for: OSINT Team, online magazine]

Mar 3 109

In System Weakness by Berend Watchus

Your Satellite Phone Is an Open Book

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. March 13, 2026 [Publication for: SYSTEM WEAKNESS, online...

Mar 13

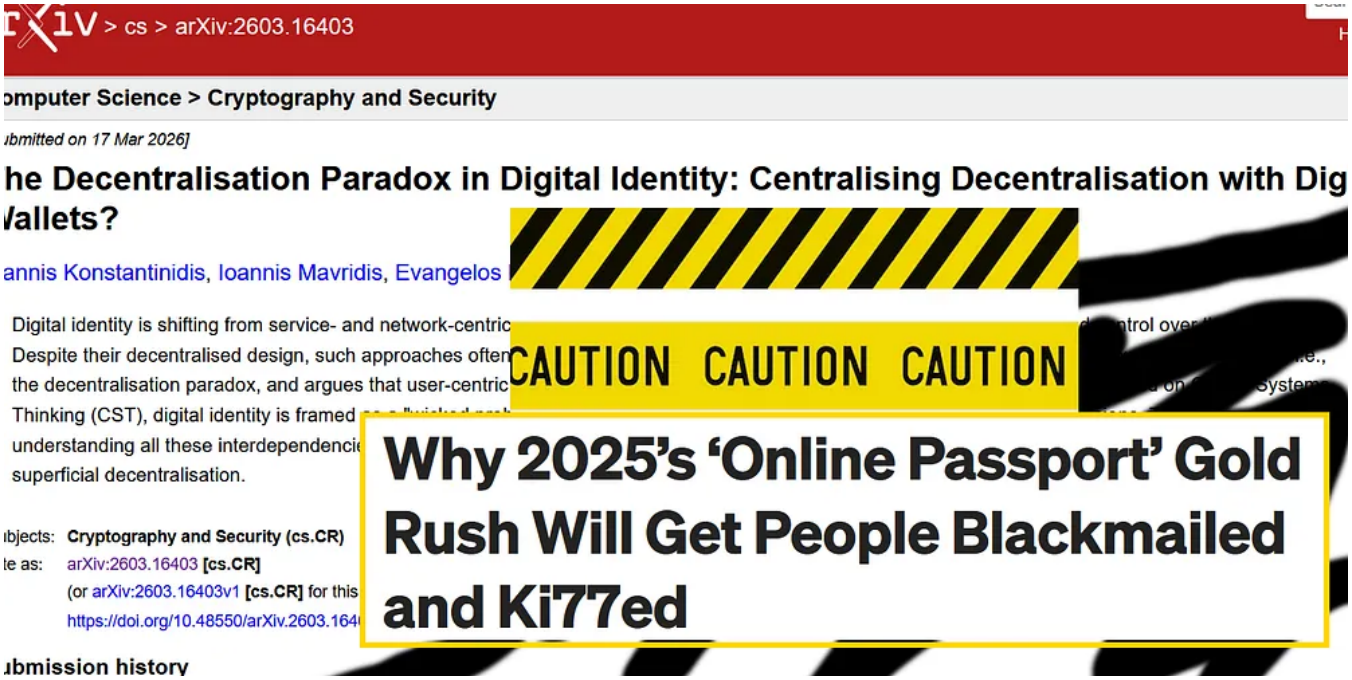

In OSINT Team by Berend Watchus

The Decentralisation Paradox Was Already Field-Tested — In Porn and webcam sexwork

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. [Publication for OSINT Team, online magazine]

2d ago 51

See all from Berend Watchus

Recommended from Medium

Will Lockett

You Have No Idea How Far Behind Tesla Is

The leapfrog moment has just happened.

6d ago 4.1K 121

In AI Software Engineer by Joe Njenga

Anthropic Just Released Claude Code Course (And I Earned My Certificate)

Anthropic just launched their Claude Code in Action course, and I’ve just passed — how about you?

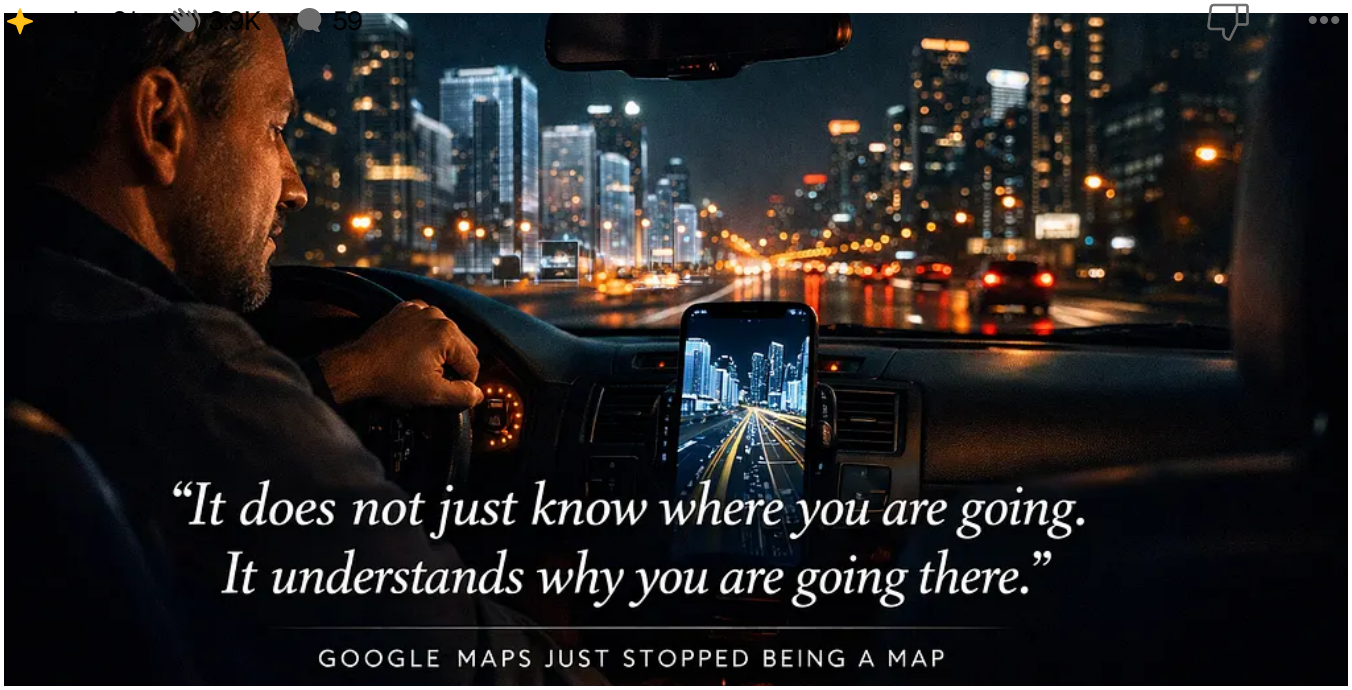

In CodeX by MayhemCode

The Biggest Google Maps Update in a Decade Just Changed How 2 Billion People Navigate

There is a version of this story where the headline is “Google Updates Its Maps App” and you scroll past it without blinking. Updates...

5d ago 41 5

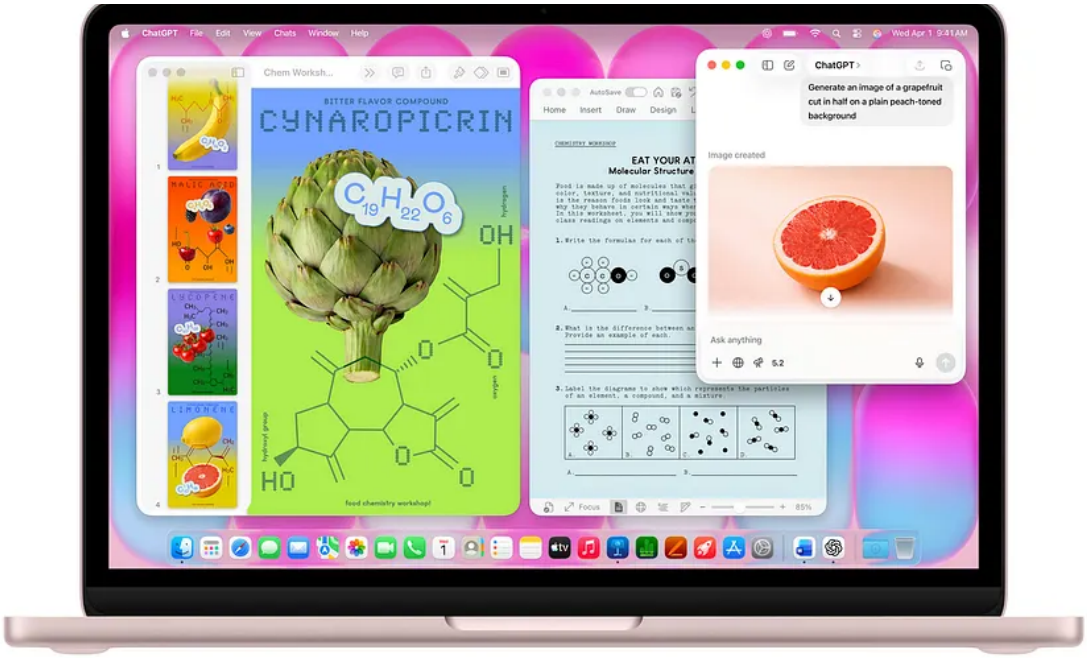

Burk

Do NOT Buy the MacBook Neo

One significant reason NOT to

5d ago 344 15

Can Artuc

VS Code Just Gave AI Full Control of Your Machine. Then Told You Not to Trust It.

VS Code 1.111 shipped a feature so dangerous that Microsoft wrote a warning against it in the same release notes. Then enabled it by...

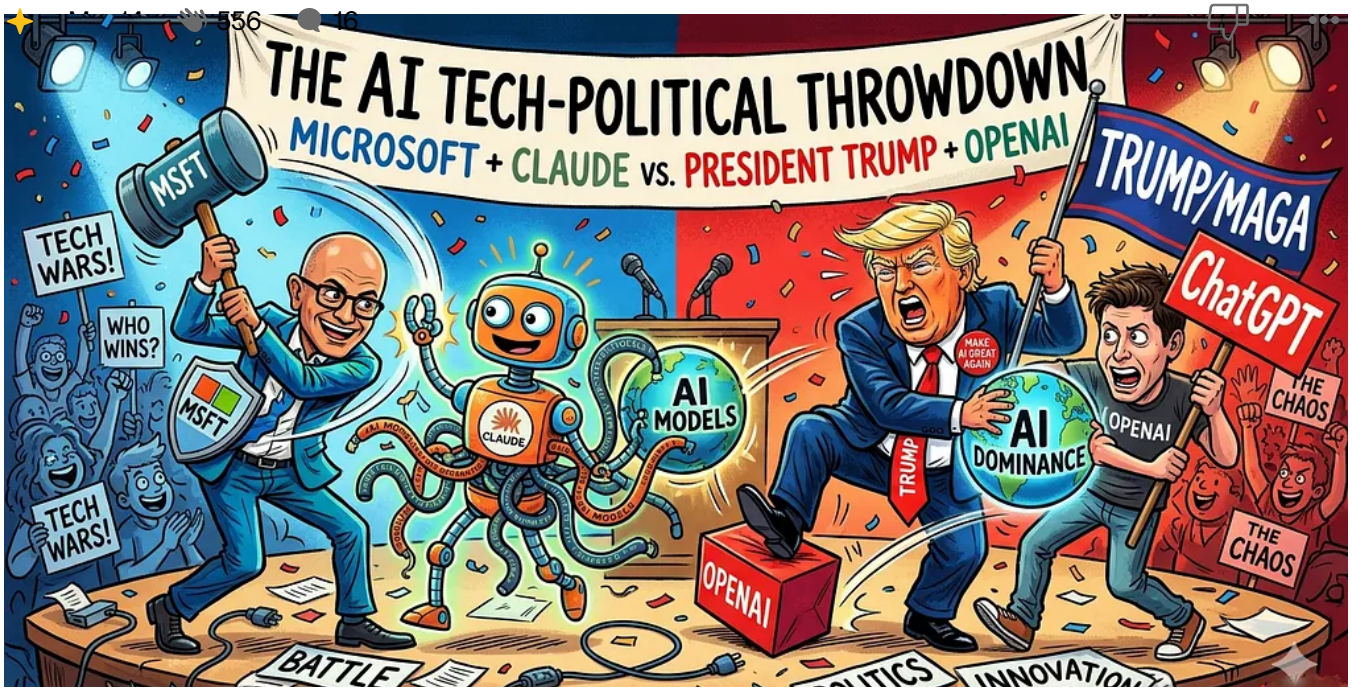

In GitBit by John Gruber

Microsoft Just Picked A Fight With One Of Its Largest Customers, President Trump

When a company known for corporate caution decides to swing at the Pentagon, something deeper is happening beneath the surface.

5d ago 169 3

See more recommendations